

IterMask²: Iterative Unsupervised Anomaly Segmentation via Spatial and Frequency Masking for Brain Lesions in MRI

Ziyun Liang¹, Xiaoqing Guo¹, J. Alison Noble¹, Konstantinos Kamnitsas^{1,2,3}

¹Department of Engineering Science, University of Oxford, Oxford, UK

²Department of Computing, Imperial College London, London, UK

³School of Computer Science, University of Birmingham, Birmingham, UK
`ziyun.liang@eng.ox.ac.uk`

Abstract. Unsupervised anomaly segmentation approaches to pathology segmentation train a model on images of healthy subjects, that they define as the ‘normal’ data distribution. At inference, they aim to segment any pathologies in new images as ‘anomalies’, as they exhibit patterns that deviate from those in ‘normal’ training data. Prevailing methods follow the ‘corrupt-and-reconstruct’ paradigm. They intentionally corrupt an input image, reconstruct it to follow the learned ‘normal’ distribution, and subsequently segment anomalies based on reconstruction error. Corrupting an input image, however, inevitably leads to suboptimal reconstruction even of normal regions, causing false positives. To alleviate this, we propose a novel iterative spatial mask-refining strategy IterMask². We iteratively mask areas of the image, reconstruct them, and update the mask based on reconstruction error. This iterative process progressively adds information about areas that are confidently normal as per the model. The increasing content guides reconstruction of nearby masked areas, improving reconstruction of normal tissue under these areas, reducing false positives. We also use high-frequency image content as an auxiliary input to provide additional structural information for masked areas. This further improves reconstruction error of normal in comparison to anomalous areas, facilitating segmentation of the latter. We conduct experiments on several brain lesion datasets and demonstrate effectiveness of our method. Code is available at: <https://github.com/ZiyunLiang/IterMask2>

Keywords: Unsupervised · Anomaly Segmentation · Brain Lesions

1 Introduction

Anomaly segmentation aims to identify patterns that deviate from a ‘normal’ distribution, defined as the distribution of training data. Such algorithms can be used for unsupervised pathology segmentation, where ‘normal’ (training) data are images of healthy tissue, while ‘anomaly’ refers to any pathology. These segmentation methods are termed unsupervised because they do not require any

‘anomalous’ images with pathology nor their manual segmentations for training - they are trained using solely ‘normal’ images (without pathologies). They can then be used to segment *any* anomalies (pathologies) in new data. This offers an alternative to the predominant supervised learning paradigm. In the context of brain lesion segmentation, the focus of this study, the common approach is supervised training of a model using manual segmentations for a *specific* pathology [10,9,8,20]. These models cannot segment other types of pathology. Moreover, it is challenging to train supervised models for rare diseases due to the scarcity of training data, and they are impractical for workflows where the type of pathology that may be in a scan is unknown beforehand. Unsupervised anomaly segmentation can instead enable development of models that can identify *any* pathology and enable applications such as automatic screening for incidental findings [19].

Related Work: The most prevalent unsupervised anomaly segmentation methods are reconstruction-based. During training, some distortion is applied to the input image and the model learns the ‘normal’ data distribution by learning to reconstruct the image. During inference, anomalies that are not seen during training are challenging to reconstruct, resulting in high reconstruction errors that are used to segment anomalies. For precise anomaly segmentation, *reconstruction errors need to be high on anomalous and low on normal areas*.

Reconstruction-based methods that leverage AutoEncoder (AE) [1,4] or Variational AutoEncoder (VAE) [4,17,22] use compression as distortion and train a model to reconstruct the image. Another type of reconstruction-based method follows the ‘corrupt and reconstruct’ paradigm: the input is corrupted during *inference*, usually by adding a type of noise, and the model is tasked to reconstruct it. Such state-of-the-art methods use diffusion models [18,5], which corrupt inputs with Gaussian noise. In both above approaches, increasing input distortion (compression or noise) amplifies reconstruction error of anomalies, since the model has not learned their appearance during training and fails to reconstruct them. This however poses a *sensitivity-precision trade-off*, because more distortion also increases reconstruction error of normal areas, causing false positives [6]. [5] proposes to address this trade-off by first generating a mask that covers all anomalies, and then by iteratively reconstructing the masked area with a diffusion model. This relates to our work that also uses an iterative process, but differs conceptually. In that method, the initial mask, which can be suboptimal, is not refined throughout iterations. Furthermore, their generative model does not have any auxiliary guidance for reconstructing the masked content, which hinders faithful reconstruction of normal tissues beneath the mask. Our method effectively addresses these points as we will demonstrate. Recently, cross-modality translation has been proposed as an intermediate step for reconstruction [12]. A model is trained for translating normal tissues between MRI modalities. It is assumed that during inference the model will fail to translate anomalies, enabling their segmentation. This approach, however, requires multi-modal data. Finally, promising results are achieved by denoising methods [11]. Unlike previously discussed approaches, the input is corrupted by adding noise only during *training* (not inference) and a model is trained to remove it. During

inference, it is assumed anomalies will be perceived as noise and will be removed. For this, noise added during training must model the appearance of anomalies and thus requires prior knowledge of their appearance. This is impractical when developing a model for segmenting ‘any’ pathology as herein.

The sensitivity-precision trade-off associated with input corruption in previous works motivates our study: Can we leverage input corruption to **amplify anomaly reconstruction error** for its segmentation, while **avoiding information loss in normal areas** to improve their reconstruction and reduce false positives, **without prior knowledge** about appearance of anomaly?

Contributions: This work introduces IterMask², a novel iterative unsupervised anomaly segmentation method that integrates spatial and frequency masking. We distort the input using spatial masking and reconstruct masked areas via a model trained solely on normal data. We **iteratively refine the spatial masks**, uncovering areas where low reconstruction error indicates normality with high confidence. This gradually re-introduces information about normal areas to further improve the reconstruction of neighbouring normal areas, reduces false positives, and shrinks the mask towards containing solely the anomaly. To further guide the reconstruction of masked normal areas, we also propose **providing structural information in the form of high-frequency image components (low-frequency masking)** as additional input. This facilitates more faithful reconstruction of normal areas, whose appearance is predictable from training, in contrast to anomalies whose structure and appearance are unexpected. We conducted extensive experiments on several brain lesion datasets that demonstrate the effectiveness of our method.

2 Method

Overview of our method is shown in Fig. 1. The method follows the ‘corrupt and reconstruct’ paradigm. To mitigate the sensitivity-precision trade-off of this type of method (Sec. 1), we propose an *iterative spatial mask refinement process* that uses spatial masking as corruption to amplify reconstruction errors on anomalies while minimizing information loss in normal areas by iteratively shrinking a spatial mask toward the anomaly. The mask shrinking process gradually introduces more information about normal-tissue appearance to the model (Fig. 1(c)). In each iteration, the spatially unmasked (ideally normal) area guides the reconstruction of the spatially masked area to shrink towards anomalies. Areas where the error map indicates confidence are considered normal and are removed from the mask for next iterations. Given the vast diversity of human brain structures, however, reconstructing masked areas without guidance would likely lead to significant reconstruction errors even for normal tissue areas, causing false positives. Therefore, we further propose extracting structural information from the image and using it as auxiliary input to the network to guide the reconstruction of masked areas. The structural information is *high-frequency image components from the Fourier domain*, as shown in Fig.1(a). The *training process* is shown in Fig.1(b). Below we describe these components in detail.




Fig. 1: Overview of proposed method. (a) The frequency masking strategy that provides structural information for an image via its high frequency components. (b) Training of the reconstruction model. Random spatial masking is applied to input image $\mathbf{x}$. High image frequencies from step (a) are given as auxiliary input $\mathbf{x}_c$ to the model. (c) shows the iterative mask refinement process during inference. The mask $\mathbf{m}_t$ of iteration t gradually shrinks toward the anomaly with the guidance of spatially unmasked area $\mathbf{x}_m$ and high image frequencies $\mathbf{x}_c$.

Iterative Spatial Mask Refinement for Inference During inference, the model aims to segment any pathology in a test dataset of brain scans $D_{te} = \{\mathbf{x}^i\}_{i=1}^N$. Assuming that the model was trained on data containing only healthy tissues, which define the ‘normal’ distribution, at inference pathologies are to be segmented as ‘anomalous’ areas with patterns deviating from the normal data distribution. For this, we use model f that reconstructs spatially masked image areas. Given an input image $\mathbf{x}$ that we aim to segment, an initial mask $\mathbf{m}_1$ is created that covers the entire brain. We will shrink the mask iteratively, with $\mathbf{m}_t$ being the mask at iteration t . The masked area is covered with Gaussian noise ϵ , while the unmasked area is copied from the input image. Therefore, the input to the model is $\hat{\mathbf{x}}_t = \mathbf{m}_t \times \epsilon + (1 - \mathbf{m}_t) \times \mathbf{x}$. The model reconstructs the image giving output $\mathbf{x}'_t = f(\hat{\mathbf{x}}_t, \mathbf{x}_f)$. Here, auxiliary input $\mathbf{x}_f$ is high-frequency image content, described in the next part of this section. We then compute the error map between $\mathbf{x}'_t$ and original image $\mathbf{x}$ via L_2 distance $\mathbf{e}_t = \|\mathbf{x}'_t - \mathbf{x}\|_2$. Since we are constantly shrinking the mask towards anomalies, we only consider the error under masked area, $\mathbf{e}'_t = \mathbf{e}_t \times \mathbf{m}_t$. To determine how to shrink the mask for the next iteration, we compare $\mathbf{e}'_t$ with threshold τ . Pixels below τ are considered confidently normal as per the model and are removed from the mask. A detailed analysis on τ can be found in Section 3. The first iteration is treated as a special case for convenience. Mask $\mathbf{m}_2$ for the second iteration is calculated using the 40th percentile of the first iteration’s error map $\mathbf{e}'_1$ as threshold τ , unmasking

40% of the brain as healthy, for faster mask convergence. The iterative process terminates automatically when the mask shrinks less than 1%.

Frequency Masking: Frequency domain analysis is fundamental in medical imaging, such as for MRI reconstruction, image denoising and more. Low frequencies of an image’s Fourier spectrum capture the mean image intensity (DC signal) and intensities of large image components. High frequencies, on the other hand, capture fine details like edges, boundaries between tissues, and fine outlines of structures [7]. Amplitude captures low-level statistics while phase is necessary for restructuring the Fourier components back to the original image [16,15].

Our method aims to reconstruct areas of an image masked in the spatial domain. To guide the reconstruction process of masked areas and enable more faithful reconstruction of normal tissues, we provide structural image information to the model in the form of high-frequency and phase image components. This reduces false positives by lowering reconstruction error of normal tissue without accompanied decrease for anomalous areas, because the model has not learned how to recreate appearance of anomalies from their high frequency components.

Given a 2D image $\mathbf{x}$ (slice of 3D scan) with size $H \times W$, we use the Discrete Fourier Transform $\mathcal{F}$ to map the input image to the frequency domain, obtaining: $\mathcal{F}(\mathbf{x})(a, b) = \sum_{h=0}^{H-1} \sum_{w=0}^{W-1} e^{-i2\pi(\frac{ah}{H} + \frac{bw}{W})} \mathbf{x}(h, w)$, where $\mathbf{x}(h, w)$ is the spatial pixel value, and the output is complex frequency value. Amplitude and phase are computed by $\mathcal{A} = |\mathcal{F}(\mathbf{x})(a, b)|$ and $\phi = \tan^{-1} \frac{\text{Im}(\mathcal{F}(\mathbf{x})(a, b))}{\text{Re}(\mathcal{F}(\mathbf{x})(a, b))}$, where Im and Re refers to imaginary and real of the complex value respectively. To capture high-level structural information, we use a high-frequency filter to mask out low-frequency amplitude while keeping all the phase information. Our high-pass filter for amplitude masking is:

$$N(u, v) = \begin{cases} 1, & \text{if } d((c_u, c_v), (u, v)) > r \\ 0, & \text{else} \end{cases}, \quad \mathcal{A}_m = \mathcal{A} \times N \quad (1)$$

Finally, the masked amplitude and phase are combined by $\mathcal{A}_m * e^{j*\phi}$ before inverse Fourier transform $\mathbf{x}_f = \mathcal{F}^{-1}(\mathbf{x})$ map it back to image space.

Spatial Random Masking for Training: To realize the aforementioned process, we train a UNet as model f , which takes spatially masked images from the training set $D_{tr} = \{\mathbf{x}^i\}_{i=1}^N$ as input and outputs the reconstructed result, with the high image frequencies x_f as auxiliary input. As there are no anomalies in training, we generate randomly shaped masks for training. For this, we first select a random point (μ_x, μ_y) in brain area and generate a multivariate 2D Gaussian distribution with the selected point as center. The probabilistic density function of the Gaussian distribution is $p(x, y) = \frac{1}{2\pi\sigma^2} e^{-[(x-\mu_x)^2 + (y-\mu_y)^2]/(2\sigma^2)}$, where (x, y) is the position. We sample the covariance σ for generating a Gaussian distribution with a dense region that won’t cover the whole brain. Then, the probabilistic density function $p(x, y)$ is used as the probability to sample masked

Table 1: Performance for different anomaly segmentation methods on all modalities of BraTS. Grey-shaded cells use the best threshold per-image. Best in bold.

Modality	FLAIR				T1CE				T2				T1			
Metrics	DSC	Sens.	Prec.	SSIM	DSC	Sens.	Prec.	SSIM	DSC	Sens.	Prec.	SSIM	DSC	Sens.	Prec.	SSIM
AE[4]	33.4	54.0	27.0	0.941	32.3	50.1	26.4	0.957	30.2	62.8	21.1	0.937	28.5	94.9	17.4	0.953
DDPM[18]	60.7	57.8	69.8	0.980	37.9	35.2	46.2	0.984	36.4	33.6	44.2	0.984	29.4	32.1	31.5	0.977
AutoDDPM[5]	55.5	57.5	58.7	0.937	36.9	58.9	28.9	0.903	29.7	53.0	22.3	0.907	33.5	61.1	24.7	0.910
Cycl.UNet[12]	65.0	63.4	73.9	0.864	42.6	47.5	42.9	0.923	49.5	48.8	53.4	0.831	37.0	45.8	35.2	0.926
DAE[11] [0, +∞]	79.7	79.1	84.5	0.926	36.7	42.0	36.2	0.844	69.6	68.1	75.3	0.903	29.5	61.2	20.5	0.860
DAE [−∞, 0]	28.5	94.9	17.3	0.875	34.7	37.9	39.3	0.966	28.5	94.9	17.4	0.885	47.9	53.7	50.9	0.958
DAE[−∞, +∞]	73.7	72.9	80.5	0.926	46.3	47.9	51.5	0.969	60.4	58.5	69.1	0.897	44.5	48.0	47.5	0.964
IterMask ²	80.2	81.3	83.3	0.985	61.7	59.1	70.9	0.997	71.2	74.4	72.9	0.994	58.5	56.6	67.6	0.996
DAE[11] [0, +∞]	84.2	83.1	87.0	0.926	42.1	52.5	42.8	0.844	74.5	72.1	80.1	0.903	35.3	64.3	28.3	0.860
IterMask ²	85.9	85.7	87.4	0.991	67.7	66.4	76.4	0.997	78.3	75.6	84.0	0.995	64.9	65.2	71.4	0.995

patches from the image. The patches have random side lengths $l = 4, 8, 16$, and the probability of a patch sampled at position (x, y) is $\frac{p(x, y)}{\sum_{x, y=1: X, Y} p(x, y)}$. After sampling around 1000 patches, random-shaped masks $\mathbf{m}_r$ are generated as shown in Fig.1(b). For input image $\mathbf{x}$, masked input to the model is $\hat{\mathbf{x}} = \epsilon \times \mathbf{m}_r + \mathbf{x} \times (1 - \mathbf{m}_r)$, where ϵ is Gaussian noise. The model’s output is reconstruction $x' = f(\hat{\mathbf{x}}, x_f)$. The training loss is the reconstruction error $L = \|\mathbf{x}' - \mathbf{x}\|_2$.

Note that we treat as a special case the first iteration of the iterative process described previously in this section, because the whole image $\mathbf{x}$ is masked for $t = 1$. For this, we train a second UNet separately. This model receives as input only the high-frequency image $\mathbf{x}_f$ and is trained to reconstruct whole input $\mathbf{x}$.

3 Experiment

Setup: We evaluate our method by training our model on healthy (i.e. ‘normal’) 2D slices and segment brain pathologies at test time. We use the **BraTS2021** dataset [3,2] with 1251 tumor cases, **ATLAS** v2.0 [13] with 655 stroke lesion cases, and SISS subtask of **ISLES** 2015 [14] with 28 stroke lesions. For BraTS, we use the same data pre-processing and selection process as described in [12]. For ATLAS and ISLES, we perform skull-strip and apply z-score intensity normalization as pre-processing. All lesion labels are merged into a single class.

Results: To evaluate segmentation quality, we use the dice coefficient (DSC), sensitivity, and precision. To evaluate the reconstruction performance in ‘normal’ areas, we measure the structural similarity by computing the SSIM score between the original image and the reconstructed image after excluding the anomaly regions. We compare our method to a variety of baseline methods, including the autoencoder-based method (AE) [4], diffusion model-based method (DDPM) [18], diffusion model with iterative resampling AutoDDPM [5], modality translation Cyclic-UNet [12], and denoising autoencoder DAE [18]. In Cyclic-UNet, the middle modalities used are T1 (FLAIR as input) and FLAIR (otherwise).

We first evaluated performance for segmenting tumors as anomalies across all modalities of the **BraTS** dataset. As shown in white-shaded parts of Tab.1,

Table 2: Performance on ISLES and ATLAS datasets. Best marked in bold.

Modality	ISLES FLAIR				ATLAS T1			
	DSC	Sens.	Prec.	SSIM	DSC	Sens.	Prec.	SSIM
AE[4]	21.5	0.475	0.167	0.791	11.9	40.7	8.6	0.954
DDPM[18]	44.0	50.3	43.9	0.937	20.2	22.6	26.5	0.987
AutoDDPM[5]	41.1	52.6	42.9	0.862	12.7	47.4	9.2	0.892
Cycl.UNet[12]	53.4	50.5	59.4	0.842	-	-	-	-
DAE[11]	51.8	61.9	52.5	0.914	11.1	33.4	32.7	0.995
IterMask ²	55.1	59.8	58.0	0.964	35.3	34.3	41.7	0.996
DAE[11]	56.5	66.8	55.6	0.905	11.1	33.4	32.7	0.995
IterMask ²	59.6	69.1	61.7	0.956	47.5	57.7	54.2	0.989

Table 3: Ablation study. $DS_{S,M,L}$ represents Dice score of small, medium and large lesions.

Modality	BraTS FLAIR					
	DSC	Sens.	Prec.	DS _S	DS _M	DS _L
-freq	64.2	69.4	63.6	20.7	63.4	72.1
-spat	75.1	73.8	79.6	46.6	77.9	77.0
DAE[11]	79.7	79.1	84.5	47.3	81.5	83.0
IterMask ²	80.2	81.3	83.3	55.6	80.8	83.6
DAE[11]	84.2	83.1	87.0	60.6	85.3	87.0
IterMask ²	85.9	85.7	87.4	65.7	86.3	88.7

our method exhibits promising performance compared to existing methods that do not include prior information during training (i.e. AE, DDPM, AutoDDPM, Cycl.-UNet), improve dice by 15.2%, 19.1%, 21.7%, 21.5% in FLAIR, T1CE, T2, and T1 modalities respectively. DAE (shown in the $[0, +\infty]$ row shaded white in Tab.1) demonstrates promising performance in FLAIR and T2 modalities, where tumors are hyper-intense, but fails to outperform other baselines methods in T1CE and T1. This is due to hyper-intense noise being used as prior during training (i.e. by learning to remove noise of range $[0, +\infty]$). We further tested DAE with different intensity priors by adjusting the intensity of added noise. In our method, without introducing prior intensity information, we outperform DAE with all kinds of noise in the dice score. More specifically, we compared our method to DAE when the intensity prior matches with the anomaly ($[-\infty, 0]$ range of noise for T1 and $[-\infty, +\infty]$ for T1CE). We improved dice by 0.5% and 1.6% in FLAIR and T2 with hyper-intense noise, and 15.4% and 10.6% in T1CE and T1, where the noise intensity is similar to normal tissues. In terms of reconstruction performance, our method improves SSIM to 0.985 in FLAIR, 0.997 in T1CE, 0.994 in T2, and 0.996 in T1, showing our iterative mask-shrinking approach can reconstruct normal areas better than other existing methods. Fig.2 shows the qualitative result of our proposed method compared to baselines.

We further evaluated our method on stroke lesion segmentation in the ISLES and ATLAS datasets, as shown in the white shaded parts in Tab. 2. Our method shows superior performance for almost all the evaluated metrics compared to baselines, demonstrating the generalizability of IterMask² to different anomalies.

Ablation Study: We assess the effectiveness of the iterative spatial mask refinement and frequency-masking as in Tab.3. Respectively, we removed frequency guidance after the initial step (*‘freq’* row) and we used high image frequencies as input to generate a single-step prediction (removing the iterative spatial mask refinement process) (*‘spat’* row). Dice score drops by 16.0% and 5.1% respectively, showing the effectiveness of both components.

Sensitivity Analysis: We first analyze the sensitivity of our model to the **mask shrinking threshold** τ in the iterative spatial mask refinement process. We choose τ from the healthy validation set’s error map obtained by applying random masks in the same way as training. As shown in Fig. 3, we experimented with τ ranging from 60-100 percentiles of the error map’s masked area on all




Fig. 2: Visualization result for FLAIR(BraTS).




Fig. 3: Sensitivity Analysis.

modalities of BraTS. Observe that our model demonstrates stability across a broad spectrum of τ , showcasing a notable level of fault tolerance. This is credited to our way of computing the final segmentation, where we use the last iteration’s reconstruction rather than directly using the final mask. So the algorithm has some degree of fault tolerance when the mask is larger than the anomaly. In Tab.1 2 3, τ is set at 80 percentile of the error map in BraTS and ISLES, and 70 in ATLAS, values found adequate in preliminary experiments prior to this analysis (thus not over-optimized). Furthermore, we analyzed sensitivity of our algorithm to the frequency masking **radius** r . In the experiment, we set high image frequencies as input to reconstruct images in a single step, showing their influence on the reconstruction process. Observe from Fig.3 that the model’s performance is stable to radius from 10-25 in all the modalities of BraTS. In all our reported experiments, we set the radius to 15, similar to the setting in [21].

Human-AI collaboration: The threshold τ is an adjustable parameter, providing clinicians with the flexibility to interactively adjust the segmentation if the result using the optimized threshold gives unsatisfactory results. Moreover, our model’s mask-shrinking process offers an intuitive interface for human interaction by showing the intermediate decision process. Therefore, we tested our model by incrementally raising τ value during the shrinking process and selected the iteration with the optimal performance per image for evaluation, obtaining improved performance, as shown in gray-shaded cells of Tab.1 2 3. We further compared our models’ performance with the best baseline method’s optimal performance (DAE), also optimizing the threshold per image when choosing the final segmentation. We found that our model constantly achieves better performance on all datasets, revealing the promising potential to interact with clinicians. Furthermore, we categorize the tumor size into small ($size < 200$), medium ($200 \leq size \leq 800$), and large ($size > 800$) and report the performance of each subgroup, demonstrating human collaboration has the potential to notably improve performance in detecting small anomalies(as shown in gray-shaded units in Tab.3). When the anomaly is small, the model can prompt clinician intervention, potentially improving prediction accuracy with minimal human effort.

4 Conclusion

This paper introduces a novel unsupervised anomaly segmentation algorithm IterMask² for segmentation of brain lesions in MRI. The process uses iterative spatial mask refinement and is guided by structural information from high-frequency image components. Extensive experiments on brain tumor and stroke lesion datasets demonstrated the effectiveness of the method in comparison to previous works. Given the very promising results, future work could seek to extend the framework for processing 3D images instead of slices processed herein.

Acknowledgments. ZL is supported by a scholarship provided by the EPSRC Doctoral Training Partnerships programme [EP/W524311/1]. The authors acknowledge UKRI grant reference [EP/X040186/1] and EPSRC grant [EP/T028572/1]. The authors also acknowledge the use of the University of Oxford Advanced Research Computing (ARC) facility in carrying out this work(<http://dx.doi.org/10.5281/zenodo.225558>).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Atlason, H.E., Love, A., Sigurdsson, S., Gudnason, V., Ellingsen, L.M.: Unsupervised brain lesion segmentation from mri using a convolutional autoencoder. In: Medical Imaging 2019: Image Processing. vol. 10949, pp. 372–378. SPIE (2019)
2. Baid, U., Ghodasara, S., Mohan, S., Bilello, M., Calabrese, E., Colak, E., Farahani, K., Kalpathy-Cramer, J., Kitamura, F.C., Pati, S., et al.: The rsna-asnr-miccai brats 2021 benchmark on brain tumor segmentation and radiogenomic classification. arXiv preprint arXiv:2107.02314 (2021)
3. Bakas, S., Akbari, H., Sotiras, A., Bilello, M., Rozycki, M., Kirby, J.S., Freymann, J.B., Farahani, K., Davatzikos, C.: Advancing the cancer genome atlas glioma mri collections with expert segmentation labels and radiomic features. *Scientific data* **4**(1), 1–13 (2017)
4. Baur, C., Denner, S., Wiestler, B., Navab, N., Albarqouni, S.: Autoencoders for unsupervised anomaly segmentation in brain mr images: a comparative study. *Medical Image Analysis* **69**, 101952 (2021)
5. Bercea, C.I., Neumayr, M., Rueckert, D., Schnabel, J.A.: Mask, stitch, and re-sample: Enhancing robustness and generalizability in anomaly detection through automatic diffusion models. arXiv preprint arXiv:2305.19643 (2023)
6. Bercea, C.I., Rueckert, D., Schnabel, J.A.: What do aes learn? challenging common assumptions in unsupervised anomaly detection. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 304–314. Springer (2023)
7. Chappell, M.: Principles of Medical Imaging for Engineers. Springer International Publishing (2019)
8. Havai, M., Davy, A., Warde-Farley, D., Biard, A., Courville, A., Bengio, Y., Pal, C., Jodoin, P.M., Larochelle, H.: Brain tumor segmentation with deep neural networks. *Medical image analysis* **35**, 18–31 (2017)

9. Isensee, F., Jäger, P.F., Full, P.M., Vollmuth, P., Maier-Hein, K.H.: nnu-net for brain tumor segmentation. In: *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries: 6th International Workshop, BrainLes 2020, Held in Conjunction with MICCAI 2020, Lima, Peru, October 4, 2020, Revised Selected Papers, Part II* 6. pp. 118–132. Springer (2021)
10. Kamnitsas, K., Ledig, C., Newcombe, V.F., Simpson, J.P., Kane, A.D., Menon, D.K., Rueckert, D., Glocker, B.: Efficient multi-scale 3d cnn with fully connected crf for accurate brain lesion segmentation. *Medical image analysis* **36**, 61–78 (2017)
11. Kascenas, A., Pugeault, N., O’Neil, A.Q.: Denoising autoencoders for unsupervised anomaly detection in brain mri. In: *International Conference on Medical Imaging with Deep Learning*. pp. 653–664. PMLR (2022)
12. Liang, Z., Anthony, H., Wagner, F., Kamnitsas, K.: Modality cycles with masked conditional diffusion for unsupervised anomaly segmentation in mri. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 168–181. Springer (2023)
13. Liew, S.L., Lo, B.P., Donnelly, M.R., Zavaliangos-Petropulu, A., Jeong, J.N., Barisano, G., Hutton, A., Simon, J.P., Juliano, J.M., Suri, A., et al.: A large, curated, open-source stroke neuroimaging dataset to improve lesion segmentation algorithms. *Scientific data* **9**(1), 320 (2022)
14. Maier, O., Menze, B.H., Von der Gablentz, J., Häni, L., Heinrich, M.P., Liebrand, M., Winzeck, S., Basit, A., Bentley, P., Chen, L., et al.: Isles 2015-a public evaluation benchmark for ischemic stroke lesion segmentation from multispectral mri. *Medical image analysis* **35**, 250–269 (2017)
15. Oppenheim, A., Lim, J., Kopec, G., Pohlig, S.: Phase in speech and pictures. In: *ICASSP’79. IEEE International Conference on Acoustics, Speech, and Signal Processing*. vol. 4, pp. 632–637. IEEE (1979)
16. Oppenheim, A.V., Lim, J.S.: The importance of phase in signals. *Proceedings of the IEEE* **69**(5), 529–541 (1981)
17. Pawlowski, N., Lee, M.C., Rajchl, M., McDonagh, S., Ferrante, E., Kamnitsas, K., Cooke, S., Stevenson, S., Khetani, A., Newman, T., et al.: Unsupervised lesion detection in brain ct using bayesian convolutional autoencoders (2018)
18. Pinaya, W.H., Graham, M.S., Gray, R., Da Costa, P.F., Tudosiu, P.D., Wright, P., Mah, Y.H., MacKinnon, A.D., Teo, J.T., Jager, R., et al.: Fast unsupervised brain anomaly detection and segmentation with diffusion models. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 705–714. Springer (2022)
19. Rowley, P., Paukner, M., Eisenmenger, L., Field, A., Davidson, R., Johnson, S., Asthana, S., Chin, N., Prabhakaran, V., Bendlin, B., et al.: Incidental findings from 16,400 brain mri examinations of research volunteers. *American Journal of Neuroradiology* **44**(4), 417–423 (2023)
20. Stollenga, M.F., Byeon, W., Liwicki, M., Schmidhuber, J.: Parallel multi-dimensional lstm, with application to fast biomedical volumetric image segmentation. *Advances in neural information processing systems* **28** (2015)
21. Xie, J., Li, W., Zhan, X., Liu, Z., Ong, Y.S., Loy, C.C.: Masked frequency modeling for self-supervised visual pre-training. *arXiv preprint arXiv:2206.07706* (2022)
22. Zimmerer, D., Isensee, F., Petersen, J., Kohl, S., Maier-Hein, K.: Unsupervised anomaly localization using variational auto-encoders. In: *Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part IV* 22. pp. 289–297. Springer (2019)